Japanese Sentiment Classification using a Tree-Structured Long Short-Term Memory with Attention

Ryosuke Miyazaki
Graduate School of System Design
Tokyo Metropolitan University
miyazaki-ryosuke@ed.tmu.ac.jp

Mamoru Komachi
Graduate School of System Design
Tokyo Metropolitan University
komachi@tmu.ac.jp

arXiv:1704.00924v1 [cs.CL] 4 Apr 2017

Abstract

Previous approaches to training syntax-based sentiment classification models required phrase-level annotated corpora, which are not readily available in many languages other than English. Thus, we propose the use of tree-structured Long Short-Term Memory with an attention mechanism that pays attention to each subtree of the parse tree. Experimental results indicate that our model achieves the state-of-the-art performance in a Japanese sentiment classification task.

1 Introduction

Traditional approaches for sentiment classification rely on simple lexical features, such as bag-of-words, that are not effective for many sentiment classification tasks (Pang et al., 2002). For example, the sentence "Insecticides kill pests." contains both *kill* and *pests*, indicating negative polarity, but the overall expression is still deemed positive.

To address this problem of polarity shift, Nakagawa et al. (2010) presented a dependency tree-based approach for the sentiment classification of a sentence. Their method assigns sentiment polarity to each subtree as a hidden variable that is not observable in the training data. The polarity of the overall sentence is then classified by Tree-Conditional Random Field (Tree-CRF) marginalizing over the hidden variables representing the polarities of the respective subtrees. In this manner, the model can handle polarity shifting operations such as negation. However, this method suffers from feature sparseness because almost all features are combination features.

To overcome the data sparseness problem, deep neural network based methods have attracted much attention because of their ability to use dense feature representations (Socher et al., 2011, 2013; Kim, 2014; Kalchbrenner et al., 2014; Tai et al., 2015; Zhang and Komachi, 2015). In particular, tree-structured approaches called recursive neural networks (RNN) have been shown to achieve good performance in sentiment classification tasks (Socher et al., 2011, 2013; Kim, 2014; Tai et al., 2015). While Tree-CRF employs sparse and binary feature representations, RNN avoids feature sparseness by learning dense and continuous feature representations. However, annotation for each phrase is crucial for learning RNN models. However, there is no readily available phrase-level polarity annotated corpus in any language other than English, which prevents tree-structured neural models from being easily ported to other languages.

We therefore propose an RNN model with an attention mechanism to compensate for the lack of phrase-level annotation. We also augment the training example with polar dictionaries. Although Kokkinos and Potamianos (2017) also provide an attention mechanism for phrase-level annotated corpus, our model performs well on sentence-level annotated corpus through the introduction of the (1) attention mechanism and (2) polar dictionary.

The main contributions of this work are as follows:

- We show that RNN models can be learned from a sentence-level polarity-tagged corpus through the use of an attention mechanism.
- We propose augmentation of the polarity-tagged corpus through the use of polar dictionaries. RNN models can effectively take advantage of such additional resources.
- We achieve the state-of-the-art performance in a Japanese sentiment classification task.

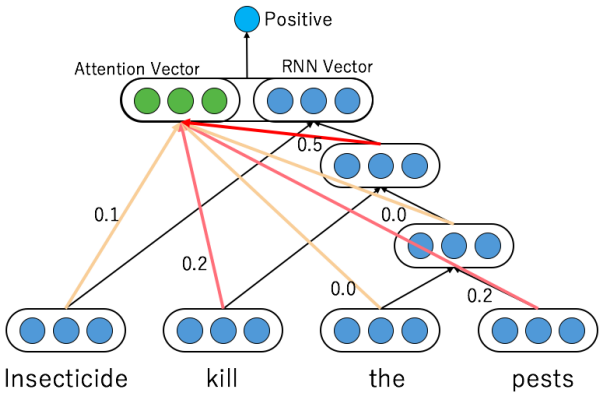

Figure 1: Sentiment classification by tree-LSTM with attention.

2 Attentional Tree-LSTM

2.1 Tree-Structured Long Short-Term Memory (LSTM)

Various RNN models for handling sentence representation considering syntactic structure have been studied (Socher et al., 2011, 2012, 2013; Qian et al., 2015; Tai et al., 2015; Zhu and Sobhani, 2015). RNN construct a sentence representation from their phrase representations by applying a composition function. Phrase representations can be calculated by recursively adopting composition functions. These RNN models are essentially identical to recurrent neural models in that they are not able to retain a long history.

Tai et al. (2015) addressed this problem by introducing LSTM (Hochreiter and Schmidhuber, 1997) to make RNN less prone to the exploding/vanishing gradient problem. In this paper, we use the Binary Tree-LSTM proposed by Tai et al. (2015) as an example of a tree-structured LSTM.

2.2 Softmax Classifier with Attention

Owing to the lack of phrase-level annotation, sentence representation may be inaccurate because it may fail to propagate errors from the root of the tree to the terminals and preterminals in a long sentence. We propose an attention mechanism to address this problem. This so-called classifier with attention takes an attention vector representation a_j in addition to a hidden representation h_j as inputs:

$$\hat{p}_\theta(y \mid h_j) = \mathrm{softmax}\left(W^{(s')} \begin{bmatrix} a_j \\ h_j \end{bmatrix} + b^{(a)}\right), \quad (1)$$

$$a_j = \sum_i a_{ji} \odot h_i, \quad (2)$$

$$a_{ji} = \frac{g(h_i, h_j)}{\sum_{i'} g(h_{i'}, h_j)}, \quad (3)$$

$$g(h_i, h_j) = \exp\left(W^{(a2)} \tanh\left(W^{(a1)} \begin{bmatrix} h_i \\ h_j \end{bmatrix}\right)\right), \quad (4)$$

where $W^{(s')} \in R^{d^l \times 2d}, W^{(a1)} \in R^{d^a \times d}$, and $W^{(a2)} \in R^{1 \times d^a}$ are the parameter matrices. In Eq.4, the biases for both $W^{(a1)}$ and $W^{(a2)}$ are omitted for simplicity. The attention vector a_j represents how much the classifier pays attention to the children nodes of the target node. The scalar values a_{ji} for each node are used to determine the attention vector. Figure 1 represents the softmax classifier with attention.

2.3 Distant Supervision with Polar Dictionaries

Unlike the Stanford Sentiment Treebank, which is annotated with phrase-level polarity, other multilingual datasets contain only sentence-level annotation. As shown in Section 3, sentiment classification without a phrase-level annotated corpus will not learn sentence representations appropriately. However, although a phrase-level polarity-tagged corpus is difficult to obtain in many languages, polar dictionaries are easy to compile (semi-)automatically. Therefore, we opt for the use of polar dictionaries as an alternative source of sentiment information.

We utilize the same polar dictionaries for short phrases and words as used in Nakagawa et al. (2010). The phrase in the training sets that matches an entry in the polar dictionaries is annotated with corresponding polarity. The key difference from Nakagawa et al. (2010) is that we use polar dictionaries as a hard label in a manner similar to distant supervision (Mintz et al., 2009); in contrast, in the previous work, it was used as a soft label for an initial hidden variable in Tree-CRF. Teng et al. (2016) also incorporated sentiment lexicons into an recurrent neural network model. Their method predicts weights for each sentiment score of subjective words to predict a sentence label. Our method uses polar dictionaries only during the training step, while the method by Teng et al. (2016) needs polar dictionaries for both training and decoding.

2.4 Learning

The cost function is a cross-entropy error function between the true class label distribution, t,(i.e.,

one hot distribution for the correct label) and the predicted label distribution, $\hat{y}$, at each labeled node:

$$J(\theta) = -\sum_{k=1}^{m} t_k \log \hat{y}_k + \frac{\lambda}{2}\|\theta\|_2^2, \qquad (5)$$

where m is the number of labeled nodes in the training set[1], and λ denotes an L2 regularization hyperparameter.

3 Experiments

We conducted sentiment classification on a Japanese corpus.

3.1 Data

We obtained pretrained word representations from word2vec[2] using skip-gram model (Mikolov et al., 2013a,b,c). We learned 200 dimensional word representations on Japanese Wikipedia's dump data (2014.11) segmented by KyTea[3] (Neubig et al., 2011). For constituency parsing, we used Ckylark (Oda et al., 2015) as of 2016.07 with KyTea for word segmentation.

We employed a Japanese polar dictionary composed by Kobayashi et al. (2005) and Higashiyama et al. (2008)[4] that contains 5,447 positive and 8,117 negative expressions. We used the NTCIR Japanese opinion corpus (NTCIR-J), which includes 997 positive and 2,400 negative sentences (Seki et al., 2007, 2008). The corpus comprised two NTCIR Japanese opinion corpora, the NTCIR-6 corpus and the NTCIR-7 corpus, as in (Nakagawa et al., 2010). We performed 10-fold cross-validation by randomly splitting each corpus into 10 parts (one for testing, one for development, and the remaining for training).

3.2 Methods

We compared our method to four baselines. All input word vectors other than those for MFS and Tree-CRF were pretrained by word2vec. We implemented our method, LogRes, and Tree-LSTM using Chainer (Tokui et al., 2015).

The following methods were used:

[1] If the data set contains only sentence-level annotation, m equals to the size of the data set.

[2] https://code.google.com/archive/p/word2vec/

[3] KyTea version-0.4.7.

[4] http://www.cl.ecei.tohoku.ac.jp/index.php?OpenResources/JapaneseSentimentPolarityDictionary

Parameter	Value
Hidden vector size	200
Optimizer	AdaDelta
Weight decay rate	0.0001
Gradient clipping	5

Table 1: The best hyperparameters.

Method	Accuracy
MFS	0.704
LogRes	0.771
Tree-CRF (Nakagawa et al., 2010)	0.826
Tree-LSTM	0.709
Tree-LSTM w/attn	0.810
Tree-LSTM w/dict	0.829
Tree-LSTM w/attn, dict	**0.844**

Table 2: Accuracy of each method on Japanese sentiment classification task.

MFS. A naïve baseline as it always selects the most frequent polarity (which is negative in this case).

LogRes. A linear classifier using logistic regression. The input features are an average of word vectors in a sentence.

Tree-CRF. A dependency-based tree-structured CRF (Nakagawa et al., 2010). It is the state-of-the-art method among our experimental datasets.

Tree-LSTM. The LSTM-based recursive neural network (Tai et al., 2015).

Tree-LSTM w/attn, dict. Our proposed method, which classifies polarity using attention and/or polar dictionaries.

3.3 Hyperparameters

We tuned hyperparameters on each development set of 10-fold cross-validation. The best parameters are shown in Table 1.

4 Results

The experimental results are shown in Table 2. The accuracy of RNN is much lower than that of the MFS baseline; moreover, Tree-LSTM, which is an improved RNN, is still lower than simple LogRes despite Tree-LSTM achieving state-of-the-art performance on the phrase-annotated Stanford Sentiment Treebank (Tai et al., 2015). In contrast, Tree-LSTM with attention achieves compa-

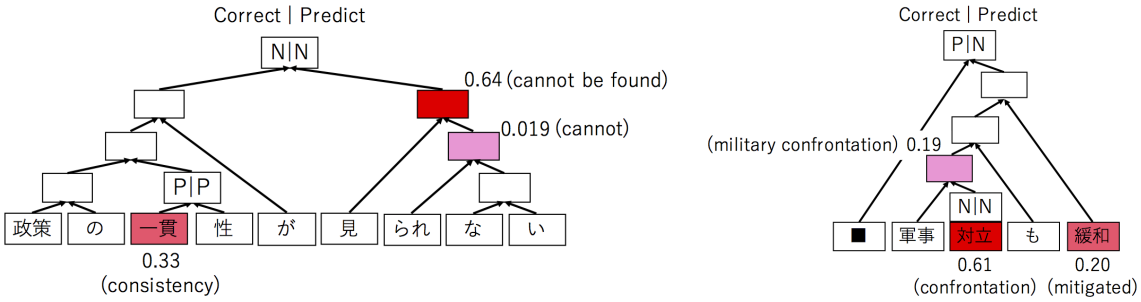

(a) Consistency of policy cannot be found.

(b) Military confrontation was mitigated.

Figure 2: Examples of our attentional Tree-LSTM sentiment classification on test set. The red square indicates a word or phrase on which great attention was paid in the softmax step, and the associated value indicates the attention weight. Root nodes are indicated by the (left)gold and (right)predicted labels ("N" indicates negative, whereas "P" indicates positive). We also show the labels for nodes that match an entry in the polar dictionary.

rable results to Tree-CRF. Our Tree-LSTM with attention and polar dictionary obtained the best accuracy.

Kokkinos and Potamianos (2017) also investigate attentional model for RNN. Their model only feeds attention vector into the softmax classifier, whereas our method inputs both attention vector and RNN vector, as illustrated in Figure 1. The accuracy of the model inputting only attention vector is 0.807, which is slightly lower than that of our proposed method.

5 Discussion

The results described above indicate that Tree-LSTM models without attention mechanism fail to learn sentence representations if phrase-level annotation is not available.

However, Tree-LSTM models can learn more accurate sentence representations if the models receive phrase-level information such as that provided by polar dictionaries. For example, in our model, attention information and polar dictionary are fed into the Tree-LSTM as phrase-level information. Although Tree-LSTM with attention and a polar dictionary outperforms Tree-CRF by 1.8 points, accuracy of Tree-LSTM without a polar dictionary is lower than that of Tree-CRF. Tree-LSTM with a polar dictionary performs better than Tree-LSTM with attention, showing that supervised label for each phrase seems to be important in learning Tree-LSTM models. Note that although our training size is 10 points lower than that in Nakagawa et al. (2010) because they used 90% of the corpus for training, our method outperforms their method in terms of accuracy.

Figures 2a shows correctly classified example. Figure 2a shows that the model classifies " (consistency)" as positive and pays 1/3 attention to it in the final classification step; however, the model correctly classifies the sentence polarity as negative by considering "ł (cannot be found)" through most of the attention.

Figures 2b displays an incorrectly classified example. In Figure 2b, the model pays attention to both " (confrontation)" and " (mitigated)"; however, it fails to predict the correct polarity of the sentence. It seems that the higher attention weight for " (confrontation)" than for " (mitigated)" influences the sentence prediction. To solve these errors, the composition function should also incorporate attention mechanism to handle polarity shifting correctly.

6 Conclusion

We presented a Tree-LSTM based recursive neural network using an attention mechanism and a polar dictionary. In this method, each phrase representation is fed into a classifier to predict the polarity of a phrase based on the phrase structures. Lexical items from the polar dictionary are used as supervised labels for each corresponding phrase or word in the same manner as distant supervision. Our experimental results demonstrated that the proposed method outperforms the previous methods.

References

Masahiko Higashiyama, Kentaro Inui, and Yuji Matsumoto. 2008. Learning sentiment of nouns from selectional preferences of verbs and adjectives. In *Proceedings of the 14th Annual Meeting of the Association for Natural Language Processing*. pages 584–587. http://www.anlp.jp/proceedings/annual_meeting/2008/pdf_dir/PA2-2.pdf.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780. http://deeplearning.cs.cmu.edu/pdfs/Hochreiter97_lstm.pdf.

Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Natural Language Processing*. pages 655–665. http://www.aclweb.org/anthology/P14-1062.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. pages 1746–1751. http://emnlp2014.org/papers/pdf/EMNLP2014181.pdf.

Nozomi Kobayashi, Kentaro Inui, Yuji Matsumoto, and Kenji Tateishi. 2005. Collecting evaluative expressions for opinion extraction. *Journal of Natural Language Processing* 12(3):203–222. https://www.jstage.jst.go.jp/article/jnlp1994/12/3/12_3_203/_pdf.

Filippos Kokkinos and Alexandros Potamianos. 2017. Structural attention neural networks for improved sentiment analysis. In *arXiv*. https://arxiv.org/pdf/1701.01811v1.pdf.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. In *Proceedings of the 1st International Conference on Learning Representations*. http://arxiv.org/pdf/1301.3781v3.pdf.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*. pages 3111–3119. https://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf.

Tomas Mikolov, Wen tau Yih, and Geoffrey Zweig. 2013c. Linguistic regularities in continuous space word representations. In *Proceedings of Human Language Technologies: The 2013 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. pages 746–751. http://www.aclweb.org/anthology/N13-1090.

Mike Mintz, Steven Bills, Rion Snow, and Dan jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. pages 1003–1011. http://web.stanford.edu/~jurafsky/mintz.pdf.

Tetsuji Nakagawa, Kentaro Inui, and Sadao Kurohashi. 2010. Dependency tree-based sentiment classification using CRFs with hidden variables. In *Proceedings of Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. pages 786–794. https://aclweb.org/anthology/N/N10/N10-1120.pdf.

Graham Neubig, Yosuke Nakata, and Shinsuke Mori. 2011. Pointwise prediction for robust, adaptable Japanese morphological analysis. In *Proceedings of The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. pages 529–533. http://www.phontron.com/paper/neubig11aclshort.pdf.

Yusuke Oda, Graham Neubig, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2015. Ckylark: A more robust PCFG-LA parser. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*. pages 41–45. http://www.aclweb.org/anthology/N15-3009.

Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*. pages 79–86. http://www.cs.cornell.edu/home/llee/papers/sentiment.pdf.

Qiao Qian, Bo Tian, Minlie Huang, Yang Liu, Xuan Zhu, and Xiaoyan Zhu. 2015. Learning tag embeddings and tag-specific composition functions in recursive neural network. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. pages 1365–1374. http://aclweb.org/anthology/P/P15/P15-1132.pdf.

Yohei Seki, David Kirk Evans, and Lun-Wei Ku. 2007. Overview of opinion analysis pilot task at NTCIR-6. In *Proceedings of the 6th NTCIR Workshop*. pages 265–278. http://research.nii.ac.jp/ntcir/ntcir-ws6/OnlineProceedings/NTCIR/81.pdf.

Yohei Seki, David Kirk Evans, and Lun-Wei Ku. 2008. Overview of multilingual opinion analysis task at NTCIR-7. In *Proceedings of the*

7th NTCIR Workshop. pages 265–278. https://pdfs.semanticscholar.org/0147/aced13e485461db1b36d7fad24adf6e0a217.pdf.

Richard Socher, Brody Huvaland Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. pages 1201–1211. http://www.aclweb.org/anthology/D12-1110.

Richard Socher, Jeffrey Pennington, Eric H. Huang, Andrew Y. Ng, and Christopher D. Manning. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. pages 151–161. http://www.anthology.aclweb.org/D/D11/D11-1014.pdf.

Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. pages 1631–1642. http://nlp.stanford.edu/~socherr/EMNLP2013_RNTN.pdf.

Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. pages 1556–1566. http://www.aclweb.org/anthology/P15-1150.

Zhiyang Teng, Duy-Tin Vo, and Yue Zhang. 2016. Context-sensitive lexicon features for neural sentiment analysis. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. pages 1629–1638. http://aclweb.org/anthology//D/D16/D16-1169.pdf.

Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. 2015. Chainer: a next-generation open source framework for deep learning. In *Proceedings of Workshop on Machine Learning Systems in The 29th Annual Conference on Neural Information Processing Systems*. http://learningsys.org/papers/LearningSys_2015_paper_33.pdf.

Peinan Zhang and Mamoru Komachi. 2015. Japanese sentiment classification with stacked denoising auto-encoder using distributed word representation. In *Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation*. pages 150–159. http://bcmi.sjtu.edu.cn/~paclic29/proceedings/PACLIC29-1018.33.pdf.

Xiaodan Zhu and Parinaz Sobhani. 2015. Long short-term memory over recursive structures. In *Proceedings of the 32nd International Conference on Machine Learning*. pages 1604–1612. http://jmlr.org/proceedings/papers/v37/zhub15.pdf.